

EV368629659

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**STRATEGY FOR SELECTIVELY MAKING NON-PUBLIC RESOURCES
OF ONE ASSEMBLY VISIBLE TO ANOTHER**

Inventors:

James H. Hogg

Christopher W. Brumme

James S. Miller

and

William G. Evans

ATTORNEY'S DOCKET NO. MS1-1872US

TECHNICAL FIELD

This subject matter relates to a strategy for compiling and executing code units, and, in a more particular implementation, to a strategy for compiling and executing code units in a virtual machine environment.

BACKGROUND

A virtual machine (VM) environment executes programs in a manner which is generally independent of the underlying complexities in the platform used to implement the execution. Microsoft Corporation's .NET Framework (provided by Microsoft Corporation of Redmond, Washington) provides a virtual machine environment with its Common Language Runtime (CLR) functionality. Sun's Java (provided by Sun Microsystems, Inc. of Santa Clara, California) provides another well known virtual machine environment.

Fig. 1 shows an overview of the operation of an exemplary architecture 100 employed by the .NET Framework in generating and executing code. Broadly stated, the architecture 100 performs a series of actions 102 used to generate intermediate language (IL) code, and another series of actions 104 for loading and executing that intermediate language code in a specific computing platform.

To begin with, source code 106 (such as program "Application.cs") is fed to a compiler 108. The source code can be written in a variety of languages that target the .NET Framework, such as C#, Visual Basic.NET, JScript.NET, C++ with Managed Extensions, and so on. The architecture 100 can use different kinds of compilers depending on the language used to create the source code 106. For instance, Fig. 1 shows the exemplary use of a C# compiler 108 to transform a source program written in C#.

The compiler 108 produces the intermediate language (IL) code and metadata 110. This IL is expressed in a common coding framework regardless of the language that

1 was used to create the original source code 106. Metadata provides descriptive
2 information regarding the IL. In general, the common coding framework provided by the
3 IL and the metadata 110 are specifically tailored to make use of the common resources of
4 the .NET Framework environment, and, in particular, to make use of the CLR. Code that
5 is built to utilize the CLR is referred to as “managed code.”

6 A key feature provided by the CLR is its use of a common set of programming
7 types via a Common Type System (CTS). Generally, a type pertains to rules used to
8 interpret information in the source code. CTS specifies a variety of types, including
9 classes, interfaces, arrays, delegates, etc. A class refers to a collection that can include
10 methods, properties, and events. An interface also refers to a collection that can include
11 methods, properties, and events; interfaces differ from classes in that they specify
12 software contracts, without implementation. An array refers to a collection of values
13 having the same type. A delegate refers to a link to a method. The above-mentioned
14 metadata generated by the compiler 108 provides descriptive information pertaining to
15 these types, such as the name of the type, the identity of any interfaces which the type
16 references, the identity of any methods implemented by the types, and so on. Additional
17 information regarding the basics of the .NET Framework can be found in a number of
18 introductory texts, such as Pratt, *Introducing Microsoft .NET*, Third Edition, Microsoft
19 Press, 2003.

20 In the execution phase 104, the architecture 100 uses a CLR loader and a just-in-
21 time (JIT) compiler (generally illustrated in Fig. 1 as a “loader/JIT” component 112) to
22 transform the IL and metadata 110 into native code 114. More specifically, the
23 loader/JIT component 112 provides code access security, cross-language integration,
24 debugging, object lifetime management, and so on. The native code 114 constitutes the
25 actual machine code that will run on an execution platform (e.g., a specific computing

1 machine). The compilation process provided by the loader/JIT component 112 is referred
2 to as “just-in-time” because the compilation can take place just prior to the execution of
3 the code.

4 The architecture 100 performs the above-identified tasks in the context of units of
5 code referred to as “assemblies.” An assembly refers to a collection of one or more files
6 that are versioned, security coded, and deployed as a unit. An assembly manifest
7 provides assembly metadata that describes various characteristics of the assembly files,
8 and which effectively ties these assembly files together. One piece of metadata contained
9 within the assembly manifest is an indication of whether the resources contained within
10 the assembly are marked as non-public resources or public resources. Resources of an
11 assembly that are designated as non-public resources are only “visible” to “entities”
12 within the defining assembly, such as by methods contained within the defining
13 assembly. On the other hand, resources of an assembly that are designated as public are
14 visible to entities both within and outside the defining assembly, such as methods
15 provided in another assembly. During execution of an assembly that references another
16 assembly’s resources, the loader/JIT component 112 reads the metadata of that other
17 assembly to determine whether these resources are public; if so, then execution can
18 proceed, and if not, execution cannot proceed.

19 Fig. 2 broadly illustrates the above-described concepts in an exemplary scenario
20 200. The scenario 200 shows four compiled assemblies, including assembly A 202,
21 assembly B 204, assembly C 206, and assembly D 208. In this exemplary case, assembly
22 A 202 includes public resources 210 that are visible to other assemblies (such as
23 assembly B 204, assembly C 206, and assembly D 208). Assembly A 202 also includes
24 non-public resources 212 that are not visible to other assemblies (e.g., assembly B 204,
25 assembly C 206, and assembly D 208). These public and non-public resources (210, 212)

1 can correspond to classes or other types within assembly A 202. Fig. 2 illustrates the
2 above-identified concepts by showing an arrow without an "X" that points from
3 assemblies B, C, and D (204, 206, 208) to the public resources 210, and an arrow with an
4 "X" that points from assemblies B, C, and D (204, 206, 208) to the non-public resources
5 212. That is, the arrow without the "X" indicates that visibility is permitted, and the
6 arrow with the "X" indicates that visibility is precluded (or blocked).

7 By virtue of the above-described designation of resources as public and non-
8 public, the loader/JIT component 112 will not allow any "outside" assembly (such as
9 assembly B 204, assembly C 206, and assembly D 208) to access assembly A 202's non-
10 public resources 212. Thus, designating the resources 212 as non-public serves a security
11 role by preventing other assemblies from tampering with potentially private and sensitive
12 information provided by the assembly A 202.

13 Nevertheless, there is a price to pay for the above-mentioned security benefits.
14 Namely, some of the outside assemblies may have a legitimate (e.g., non-malicious) need
15 to access the non-public resources 212 of assembly A 202, but the non-public status of
16 these resources 212 prevents these outside assemblies from accessing these resources
17 212. One way to address this problem is to re-designate the non-public resources 212 as
18 public resources. This opens up the previously non-public resources 212 to all legitimate
19 uses by outside assemblies; however, the public designation is not discriminatory, so that
20 it also accommodates assemblies that wish to access the resources 212 for non-legitimate
21 (e.g., malicious) uses. This can result in a security breach and potential corruption of
22 computer code and database records. Further, even where there is no security risk
23 involved in providing assembly resources to a wider population of outside assemblies, the
24 designer of the assembly may nonetheless wish to restrict the population of interacting
25 assemblies so as to simplify the assembly's design and testing; that is, an assembly may

1 need to adopt a more robust design if it is intended to interact with a wider population of
2 assemblies than originally intended.

3 Another way to permit an outside assembly to access the non-public resources
4 212 of assembly A 202 is to redesign that outside assembly such that it actually
5 incorporates all of the resources that its needs from assembly A 202. However, this
6 change is typically a relatively invasive, error prone, and time intensive task, requiring
7 redesign of source code, recompiling, and so on. Also, this solution reduces the
8 modularity and independence of a software product's computer code, which is generally
9 considered a negative consequence in the computing arts.

10 Accordingly, there is an exemplary need in the art to provide a more efficient way
11 of permitting an entity to selectively gain access to resources in an assembly that have
12 been designated as non-public.

13 14 **SUMMARY**

15 According to one exemplary implementation, a method is described for forming
16 (e.g., compiling) and executing code units. The method includes a step of forming a first
17 code unit having a resource, and forming a second code unit which references the
18 resource in the first code unit. The first code unit places restrictions on the visibility of
19 the resource to other code units, but the first code unit also includes an attribute which
20 overrides the restrictions with respect to the second code unit. The method also includes
21 a step of executing the first and second code units using a runtime component, involving
22 making the resource of the first code unit visible to the second code unit as instructed by
23 the attribute.

24 Additional exemplary implementations are also described in the following.
25

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows an exemplary .NET Framework architecture.

Fig. 2 shows an exemplary assembly having resources designated as public and resources designated as non-public, and also shows the consequences thereof to other assemblies that wish to access the non-public resources.

Fig. 3 shows an exemplary assembly having resources designated as public and resources designated as non-public, where an InternalsVisibleTo (IVT) attribute is associated with the non-public resources to designate which outside assembly is allowed to access these non-public resources.

Fig. 4 shows two applications of an exemplary source code compiler used to respectively build two of the assemblies shown in Fig. 3.

Fig. 5 shows an exemplary loader/JIT component used to execute one of the assemblies shown in Fig. 3.

Fig. 6 shows an exemplary procedure for building and executing assemblies using the mechanisms shown in Figs. 4 and 5.

Fig. 7 shows an exemplary computing environment for implementing aspects of the architecture shown in Figs. 4 and/or 5.

The same numbers are used throughout the disclosure and figures to reference like components and features. Series 100 numbers refer to features originally found in Fig. 1, series 200 numbers refer to features originally found in Fig. 2, series 300 numbers refer to features originally found in Fig. 3, and so on.

DETAILED DESCRIPTION

Broadly stated, the solution described here involves forming a first code unit having a resource, wherein the first code unit places restrictions on the visibility of the

1 resource to other code units. However, an attribute is added to the first code unit which
2 selectively overrides the restrictions with respect to a second code unit (or additional
3 other specified referencing modules). When the second code unit is executed at runtime,
4 the loader/JIT component makes the resource of the first code unit visible to the second
5 code unit, as instructed by the attribute. As used in this disclosure, the term “code unit”
6 broadly refers to any unit of code in any kind of code environment.

7 According to one exemplary implementation, the above-mentioned attribute can
8 be implemented as a custom attribute. This custom attribute can be added to the first
9 assembly’s metadata as produced by a language compiler. This attribute is referred to
10 herein as an “InternalsVisibleTo” attribute, or an “IVT” attribute for brevity. This label
11 is descriptive of the function performed by this attribute; namely, the InternalsVisibleTo
12 attribute specifies that an internal resource within the first code unit is made “visible to”
13 the second code unit. However, this label is otherwise exemplary and non-limiting.

14 Another term used in this discussion is that of “visibility-privileged assembly.” In
15 the context of this discussion, attaching an IVT attribute, that identifies a second code
16 unit, to a resource in the first code unit renders that second code unit a “visibility-
17 privileged” assembly with respect to the first code unit. Again, this term is descriptive of
18 the characteristics of an assembly that is granted visibility via the IVT attribute, but this
19 term is otherwise exemplary and non-limiting.

20 In one code environment, making a resource in a first code unit “visible” or
21 “available” to a second code unit does not necessarily mean that the second code unit can
22 utilize this resource, as the first code unit may impose other constraints (to be discussed
23 below) which restrict the second code unit’s attempt to utilize or access the resource.
24 However, in another code environment, making a resource in a first code unit “visible” or
25 “available” to a second code unit does mean that the second code unit can utilize this

1 resource. Accordingly, the terms “visible” and “available” should be construed broadly
2 herein. Considered abstractly, a code environment has the capability to impose any kind
3 of restriction that relates to a resource; a resource is made “visible” or “available” (as
4 these terms are broadly used herein) when it removes that restriction with respect to a
5 specified assembly; different technology-dependent code environments can apply this
6 mechanism to different kinds of restrictions. Stated in a different way, the terms
7 “visible” or “available” refer to any characteristic related to a resource which can be
8 restricted, and for which that restriction can be selectively overridden.

9 The above-summarized technique can be applied in many different programming
10 environments. It can be generally used in any environment that involves one code unit
11 referencing another, including stand-alone computer environments, network
12 environments, and so forth. To facilitate discussion, and to provide concrete examples,
13 this discussion will feature the application of the above-described technique to the type of
14 virtual machine environment illustrated in Fig. 1, and, more specifically, to Microsoft’s
15 .NET Framework environment. In connection with this environment, this disclosure will
16 discuss the code units in the specific context of assemblies that contain various resources
17 that conform to the .NET’s Common Type System (CTS). Again, however, the
18 techniques described herein can be applied to other programming environments besides
19 the .NET environment in which the term “code units” can take on a different
20 interpretation.

21 In the following discussion, any of the functions described herein can be
22 implemented using software, firmware (e.g., fixed logic circuitry), manual processing, or
23 a combination of these implementations. The term “logic” as used herein generally
24 represents software, firmware, or a combination of software and firmware. In the case of
25 a software implementation, the logic can represent collections of program code (and/or

1 declarative statements) that perform specified tasks when executed on a processing
2 device or devices (e.g., CPU or CPUs). The program code can be stored in one or more
3 computer readable memory devices.

4 Finally, a number of examples will be presented in this disclosure in the
5 alternative (e.g., case A or case B). In addition, this disclosure encompasses those cases
6 which combine alternatives in a single implementation (e.g., case A and case B), even
7 though this disclosure may not expressly mention these conjunctive implementations in
8 every occasion.

9 Section A of this description provides an overview of exemplary characteristics of
10 the visibility-privileged relationship. Section B describes an exemplary system for
11 implementing the visibility-privileged behavior described in Section A. Section C
12 describes an exemplary procedure for building and executing assemblies that are subject
13 to the visibility-privileged relationship. Section D provides additional details regarding
14 the characteristics of the visibility-privileged relationship. And Section E describes an
15 exemplary computer environment for implementing aspects of the procedure discussed in
16 Section C.

17 18 **A. Overview of Exemplary Characteristics of the Visibility-Privileged** 19 **Relationship**

20 Fig. 3 illustrates the use of the IVT attribute in the context of the scenario
21 introduced in Fig. 2. In the scenario 300 shown in Fig. 3, a language compiler has
22 produced a collection of assemblies, including assembly A 302, assembly B 304,
23 assembly C 306, and assembly D 308. These assemblies are merely exemplary; other
24 scenarios can involve more assemblies or fewer assemblies.
25

1 Assembly A 302 has a collection of resources 310 identified as public resources
2 and another collection of resources 312 identified as non-public resources. These
3 resources (310, 312) can pertain to any unit and kind of information that is typically
4 provided in an assembly, including modules in assemblies (for example, modules in
5 assemblies in the .NET code environment), namespaces, types, members, or other
6 information. As in the case of Fig. 2, the public designation of resources 310 means that
7 any outside assembly (such as assembly B 304, assembly C 306, and assembly D 308)
8 can see these public resources 310. The non-public designation of resources 312 means
9 that any outside assembly (such as assembly B 304, assembly C 306, and assembly D
10 308) cannot normally see these non-public resources 312.

11 However, unlike the scenario 200 shown in Fig. 2, the scenario 300 in Fig. 3
12 associates an IVT attribute 314 with the non-public resources 312. This IVT attribute
13 314 will be interpreted by a loader/JIT component so as to allow a designated outside
14 assembly or assemblies to access the non-public resources 312, while still preventing
15 other assemblies from accessing the non-public resources 312. This satisfies the
16 exemplary needs identified in the Background section of this disclosure; namely, those
17 assemblies that have a legitimate need to see the non-public resources 312 can do so by
18 virtue of the IVT attribute 314, while other assemblies that might not have a legitimate
19 need are prevented from doing so. Note that it is also possible to restrict resource
20 availability by making the resource public to all assemblies, yet restricting this
21 availability by associating a Strong Name Identity Permissions (SNIPs) with it; in
22 contrast, however, the above-described IVT attribute operates by maintaining the non-
23 public status of the resource, but selectively opening up availability to a specific named
24 assembly.

1 The IVT attribute 314 is shown in the high-level conceptual depiction of Fig. 3 as
2 being “attached” to or associated with the non-public resources 312. Generally, the IVT
3 attribute 314 can be associated with any one or more resources in the assembly A 302.
4 For instance, although not shown, the IVT attribute 314 can be associated with assembly
5 A 302 as a whole, such that it applies to all of the non-public resources contained therein.
6 Alternatively, the IVT attribute 312 can apply to individual resources within the assembly
7 A 302, such as any collection of modules within assembly A 302 (e.g., in the context of
8 the .NET programming environment), any collection of namespaces within assembly A
9 302, any collection of types within the assembly A 302 (e.g., any class types, array types,
10 pointer types, interface types, enumeration types, and so on), and any collection of
11 members within the types, and so on.

12 The IVT attribute 314 can be associated with any of the above-described
13 resources by positioning the IVT attribute 314 at an appropriate position within the
14 source code. The language compiler then compiles that source code (with IVT attributes
15 embedded therein) to produce assembly A 302. Alternatively, as will be described
16 below, IVT-related information can also be specified via a command line operation used
17 to invoke the language compiler, so that this IVT-related information does not have to be
18 entirely embedded within the source code itself (or embedded within the source code at
19 all). Further, although Fig. 3 only shows one IVT attribute 314 associated with assembly
20 A 302, this assembly A 302 can include multiple IVT attributes, which may respectively
21 identify plural visibility-privileged assemblies (that is, assemblies that are permitted to
22 see the non-public resources 312 of assembly A 302). In an alternative implementation, a
23 single IVT attribute can specify plurality visibility-privileged assemblies.

24 Finally, the IVT attribute 314 permits a named outside assembly to “see” the non-
25 public resources 312. This privilege can extend to all requesting entities (e.g., methods)

1 associated with the named outside assembly. However, in another implementation, the
2 IVT attribute 314 can grant more fine-grained privileges to entities in the named outside
3 assembly, only permitting certain entities within this named outside assembly to view the
4 non-public resources, or only permitting visibility under certain specified conditions. In
5 still further implementations, the IVT attribute 314 can grant visibility to entities on some
6 basis other than a per-assembly basis, e.g., by specifying that visibility is to be granted
7 for certain resource types, and so on. However, to facilitate discussion, the following
8 description is framed in the context of IVT attributes that grant access to entities on a per-
9 assembly basis.

10 Fig. 3 shows a high level depiction of an exemplary composition of the IVT
11 attribute 314. The IVT attribute 314 includes a first argument 316 that identifies which
12 assembly is designated as a visibility-privileged assembly. In this exemplary case, the
13 first argument 316 identifies the visibility-privileged assembly as assembly D 308. This
14 means that assembly D can “see” the non-public resources 312, whereas other assemblies
15 still cannot see the non-public resources 312. Fig. 3 illustrates this concept by showing
16 an arrow without an “X” between assembly D 308 and the non-public resources 312, and
17 by showing an arrow with an “X” (indicating no visibility) between assemblies B and C
18 (304, 306) and the non-public resources 312. As before, all assemblies (304, 306, and
19 308) can access the public resources 310 (if they exist within a particular assembly).
20 Further note that, in the case of Fig. 3, the IVT attribute 314 can target individual
21 resources within assembly A 302; in another implementation, however, the IVT attribute
22 314 can apply globally to all non-public resources within assembly A 302.

23 The visibility-privileged relationship is one-way (non-symmetric), and non-
24 transitive. More specifically, as stated above, assembly D 308 has visibility into the non-
25 public resources of assembly A 302. This can be represented using the exemplary

1 notation: A <:: D. The relationship is one-way, so that D <:: A does not hold. To
2 achieve symmetric behavior, assembly D 308 needs to declare assembly A 302 as a
3 visibility-privileged assembly by adding an IVT attribute to assembly D 308. In one
4 exemplary implementation, the relationship is also non-transitive, meaning that, if A <::
5 D and D <:: C, then A <:: C does not hold. To achieve this relationship, the user can add
6 an IVT attribute to assembly A that grants visibility to assembly C. Finally, the
7 visibility-privileged relationship is reflexive, such that A <:: A (meaning that assembly A
8 302 has visibility into its own resources.)

9 As indicated in Fig. 3, the IVT attribute 314 also includes a number of other
10 arguments 318. In one implementation, these other arguments 318 are mandatory. In
11 another implementation, one or more of the other arguments 318 may be optional
12 (including, in one case, all of the arguments).

13 The exemplary other arguments 318 shown in Fig. 3 can include a version
14 argument. The version argument provides a numeric value which identifies the version
15 of the visibility-privileged assembly, in this case, assembly D 308. The version
16 information can be expressed in the format of <major version>.<minor version>.<build
17 number>.<revision>. The other arguments 318 can also include a culture argument. The
18 culture argument provides information regarding the language and regional style
19 employed by the visibility-privileged assembly. For example, the culture argument can
20 specify "de"; this argument specifies that the named assembly must be tailored for user
21 interaction in the German language (for example). The other arguments 318 also include
22 security information. The security information can identify a public key token that is
23 used, or will be used, by the visibility-privileged assembly. Additional information
24 regarding these arguments will be provided in Section D.
25

1 The use of the combination of version, culture, and security code (e.g., public key
2 token) is beneficial because it conforms to other referencing schemes (e.g., AssemRefs)
3 used in the .NET Framework. However, these three arguments are otherwise merely
4 exemplary. Other implementations can provide additional arguments, fewer arguments,
5 and/or different arguments.

6 The IVT attribute 314 as a whole can be expressed using a variety of different
7 syntaxes depending on the developer's programming environment and other factors.
8 Four examples exemplary syntaxes that can be used are listed below:

- 9
- 10 a. [assembly:InternalsVisibleToAttribute("D",
11 PublicKeyToken="32ab4ba45e0a69a1")]
 - 12 b. [assembly:InternalsVisibleToAttribute("D", Version="1.2.3.4",
13 Culture="en-US")]
 - 14 c. [assembly:InternalsVisibleToAttribute("D",
15 PublicKeyToken="32ab4ba45e0a69a1", Culture="de")]
 - 16 d. [assembly:InternalsVisibleToAttribute("D",
17 PublicKeyToken="1234abcd56780000", Version="1.2.3.4", Culture="fr")]
- 18

19 Assume that these attributes have been added to the source code of assembly A 302 to
20 relax the non-public status assigned to the resources 312 (although, as will be discussed,
21 it is possible to convey the same information via a command line mechanism). Case (a)
22 pertains to a situation in which assembly A 302 has granted visibility-privileged status to
23 assembly D 308, provided that that assembly D 308 furnishes the specified public key
24 token of "32ab4ba45e0a69a1." Note that this declaration makes non-public types in the
25 declaring assembly A 302 visible to assembly D 308 for all versions and cultures of

assembly D 308 (since the IVT attribute does not specify a version argument or a culture argument). Case (b) pertains to a situation in which assembly A 302 has granted visibility-privileged status to assembly D 308, provided that assembly D 308 furnishes a version of "1.2.3.4" and conforms to a culture of "en-US" (i.e., English, United States). Case (c) pertains to a situation in which assembly A 302 has granted visibility-privileged status to assembly D 308, provided that Assembly D 308 furnishes the specified public key token of "32ab4ba45e0a69a1" and conforms to a culture of "de" (i.e., German). Case (d) pertains to a situation in which assembly A 302 has granted visibility-privileged status to assembly D 308, provided that that assembly D 308 furnishes the specified public key token of "1234abcd56780000," and conforms to a version of "1.2.3.4" and a culture of "fr" (i.e., French).

More specifically, the IVT attribute 314 can be implemented as a custom attribute, which is defined as a class. An exemplary definition of the IVT attribute 314 is specified as follows:

```
[AttributeUsage(AttributeTargets.Assembly, AllowMultiple=True,
Inherited=false)]

public sealed class InternalsVisibleToAttribute : Attribute {

    private string _name;

    private string _culture

    private byte[] _publicKeyToken

    private string _version

    public VisibleToAttribute(String assemblyName);

    public string Name { get; set; }

    public string Culture { get; set; }
```



```
1      public byte[] PublicKeyToken { get; set; }  
2      public string Version { get; set; }  
3  }
```

4
5 This custom attribute is added to assembly information produced by the compilation
6 process. This custom attribute can thereafter be interpreted by the loader/JIT component
7 in the runtime execution of assemblies.

8 9 **B. Exemplary System for Processing Visibility-Privileged Assemblies**

10 Fig. 4 shows architecture for producing exemplary assembly A 302 and assembly
11 D 308 shown in Fig. 3. That is, the left side of Fig. 4 shows a compilation process 402
12 for compiling source code using a compiler 404 to produce the previously mentioned
13 assembly 302. The right side of Fig. 4 shows a compilation process 406 for compiling
14 source code using the same compiler 404 to produce the previously mentioned assembly
15 308. In other words, the same compiler 404 is being used to compile the source code for
16 both assemblies A and D (302, 308). Note that Fig. 4 depicts the use of this compiler 404
17 in two different applications to clarify its operation when it is producing an assembly that
18 includes an IVT attribute (such as assembly A 302) and when it is producing an assembly
19 (such as assembly D 308) which references resources enabled by an IVT attribute in
20 another assembly.

21 Beginning with the process 402, there are at least two ways to add the IVT
22 attribute to the binary compiled assembly 302 through the use of the compiler 404. The
23 first technique is via the approach described in Section A. Namely, the designer can add
24 one or more IVT attributes to the source code used to produce the assembly A 302. Fig.
25 4 represents this technique by showing the input of a source code unit 408 having IVT

1 attributes 410 at appropriate locations in the source code. In one implementation, these
2 IVT attributes 410 can be used to govern the visibility of all of the resources (e.g., types)
3 in the assembly A 302. In another implementation, these IVT attributes 410 can be used
4 to govern the visibility of specified resources in the assembly A 302, such as only certain
5 types or members. In the latter case, the designer can associate the IVT attributes 410
6 with certain resources by positioning these IVT attributes 410 in the source code next to
7 the respective non-public resources that they modify. This implementation is merely
8 exemplary; other approaches can be used to associate the IVT attributes 410 with certain
9 resources depending on the syntactical framework in which a designer seeks to use the
10 visibility-privileged functionality described herein.

11 In a second technique, a designer can write a source code unit 412 that omits
12 some or all of the IVT attribute information, and then the developer can specify this IVT
13 information 414 in the command line information used to invoke the compiler 404. That
14 is, in one implementation, the source code unit 412 can contain no IVT attributes, or may
15 contain IVT attributes which omit certain arguments, such as the security codes (e.g.,
16 public key tokens) of the visibility-privileged assemblies. The user can add this missing
17 information in the command line used to invoke the compiler 404. The compiler 404
18 receives this information, processes it, and injects the attributes as metadata into the
19 binary assembly that it produces in compilation. Again, different command line syntaxes
20 can be used to express the IVT information depending on the programming environment
21 in which the assembly is being built.

22 The compiler 404 itself can be implemented as any kind of compiler used to
23 produce managed code, such as a C# compiler, a C++ with Managed Extensions
24 compiler, a Visual Basic.NET compiler, and so on. Fig. 4 generically identifies any
25 functionality configured to process IVT information as logic 416. This logic 416 can

1 represent software and/or firmware for performing the ascribed compilation tasks to be
2 described below. In one implementation, the logic 416 need not employ special
3 provisions for processing the IVT attributes in assembly A 302. This is possible because,
4 in one implementation, the IVT attributes are implemented as custom attributes, which
5 the above-mentioned types of compilers are already designed to handle. In another
6 implementation, the logic 416 can include specifically-tailored functionality for handling
7 the use of IVT attributes in assembly A 302. Further, as will be described below in the
8 context of the discussion of process 406, the logic 416 can include special provisions for
9 processing assembly D 308 which is granted access to the non-public resources in
10 assembly A 302 by virtue of an IVT attribute added to assembly A 302.

11 The result of the compilation performed by compiler 404 is the production of a
12 binary assembly file (or files), as well as an assembly manifest; together, these files
13 constitute assembly A 302. The assembly A 302 includes intermediate language (IL)
14 code and metadata which reference a number of resources expressed in the Common
15 Type System (CTS). The IVT attributes are embedded within assembly A 302 (e.g.,
16 within assembly A 302's metadata). These IVT attributes are interpreted by the
17 loader/JIT component when a customer decides to execute assemblies A and D (302,
18 308) (to be described below in the context of Fig. 5).

19 Now referring to the right portion of Fig. 4, the compilation of assembly D 308
20 involves compiling source code 418 using compiler 404 to produce assembly D 308. In
21 this case, the source code 418 for assembly D does not contain any IVT attributes, unless
22 assembly D, in turn, wishes to expose some or all of its resources to another assembly,
23 such as assembly A 302 (although this possible scenario is not illustrated in Fig. 4 so as
24 not to overly complicate the figure).

1 In compiling the source code 418, the logic 416 of compiler 404 can be
2 configured to verify that assembly D 308 is indeed allowed to access the non-public
3 resources of assembly A 302. If not, the logic 406 can be configured to generate an error
4 message, thus alerting the developer to this problem before the assemblies are shipped to
5 a customer. Alternatively, the compiler 404 need not perform any checking to determine
6 whether assembly D 308 can in fact access the IVT-enabled non-public resources of
7 assembly A 302. In this latter case, the developer may rely on the loader/JIT component
8 stage of the execution to detect any anomalies that arise in the accessing of the IVT-
9 enabled non-public resources of assembly A 302. In either case, the compiler 404
10 produces the binary assembly D 308 for shipment to a customer. The assembly D 308
11 includes intermediate language (IL) code and metadata which reference a number of
12 resources expressed in the Common Type System (CTS). Some of the metadata may
13 include reference information which points to non-public resources in assembly A 302.

14 Fig. 5 shows a process 502 for executing the assembly D 308 that was produced
15 in the process 406 of Fig. 4. The process 502 involves processing assembly D 308
16 through a loader/JIT component 504. The loader/JIT component 504 includes logic 506
17 configured to process IVT information in assemblies that are being referenced. In this
18 case, assembly D 308 includes resources which require access to the non-public resources
19 of assembly A 302. Normally, this access would be precluded because non-public
20 resources are earmarked for selective use within their local assembly, e.g., in this case,
21 assembly A 302. However, assembly A 302 includes the IVT attribute 314 that
22 selectively makes the non-public resources available to assembly D 308. The logic 506
23 can determine the availability of resources in a stepwise fashion by examining each of the
24 possible conditions under which these resources could be made available. For example,
25 the logic 506 can first determine whether the referenced external resources are public. If

not, the logic 506 can determine whether the referenced external resources are otherwise made visible by virtue of an IVT attribute. To determine whether a resource is made available by virtue of an IVT attribute in another assembly, the logic 506 can access and examine an in-memory “visibility-privileged assemblies” table that identifies assemblies and resources that have been made available via respective IVT attributes (and which also identifies to what assemblies these resources have been made available). The loader/JIT component 504 generates native code 508 for executing assembly D 308 if all of its resources are determined to be available; otherwise, the loader/JIT component 504 can be configured to produce an error message. The native code 508 is in a form that is tailored for execution on a specific machine/software platform used by the customer.

The above-described techniques for selectively granting access to resources are code language independent. For instance, in one exemplary implementation, the syntax of the IVT attributes is the same regardless of the language used to form the source code units. Also, at runtime, the logic 506 of the loader/JIT component 504 is configured to process assemblies that contain IVT attributes (or that reference other assemblies that contain resources made available by IVT attributes in these other assemblies) in a manner which is independent of the code languages used to form the assemblies. Accordingly, the IVT functionality provides a universal (e.g., generally applicable) technique for selectively granting access to resources in assemblies.

C. Method of Operation

Fig. 6 shows an exemplary procedure 600 for building and then executing assemblies, where at least one assembly has designated another assembly as a visibility-privileged assembly by virtue of the above-described techniques.

1 In step 602, a developer compiles source code using any type of compiler that
2 targets the CLR to produce a compiled binary assembly containing intermediate language
3 (IL) code and metadata. The assembly may contain one or more IVT attributes or may
4 include entities (e.g., methods) which reference the non-public resources of another
5 assembly which have been made available via an IVT attribute in that other assembly.
6 Step 604 indicates that this compiling step is repeated for each of the assemblies in a suite
7 of assemblies that form a particular software product or suite of software products. In
8 this compilation, the compiler can be configured to perform checking to ensure that the
9 IVT-enabled resources can be properly accessed; alternatively, the compiler can omit this
10 checking, instead relying on the loader/JIT component to catch any anomalies related to
11 the IVT functionality described herein.

12 In step 606, the developer ships (or otherwise transfers) the compiled assemblies
13 to a customer. In the context used here, “developer” and “customer” have broad
14 connotation. In the most common case, a developer refers to an entity (e.g., a company)
15 that produces computer code, and a customer can refer to an entity (e.g., an individual or
16 company) that purchases and executes this computer code. However, in one application,
17 the developer can produce code for “in house” use within its own organization.

18 In step 608, the customer receives the compiled assemblies. In step 610, the
19 customer executes the assemblies using the load/JIT component to produce native code
20 for execution on the customer’s computer platform. This final compiling can be
21 performed just prior to running the code (hence the name “just-in-time” compilation).
22 Alternatively, the customer can produce the native code in advance for running at some
23 later point in time. In any event, the loading and runtime compiling of an assembly that
24 references resources IVT-enabled resources in another assembly will utilize the
25

1 functionality 506 described in the context of Fig. 5. This functionality 506 will alert the
2 user to any errors in referencing IVT-enabled non-public resources.

4 **D. Additional Details Regarding the Visibility-Privileged Relationship**

5 This section provides additional information regarding features of the visibility-
6 privileged relationship. Namely, this section describes details regarding IVT attribute
7 arguments (such as the version and security arguments), details regarding how the IVT
8 attribute affects members defined in an assembly in various circumstances, details
9 regarding how the IVT attribute affects nested members in an assembly, and details
10 regarding how the IVT attribute affects reflection functionality provided in the .NET
11 programming environment.

12 *A.1. The Version Argument*

13 As mentioned above, the IVT attribute can include an optional version argument.
14 The version argument specifies the version of the visibility-privileged assembly that is
15 permitted access to the non-public resources of the visibility-enabling assembly. The
16 inclusion of a version attribute has different implications for the compilation process and
17 the runtime process (corresponding to the processes illustrated in Figs. 4 and 5,
18 respectively), as will be described below.

19 As to the compile-time processes shown in Fig. 4, recall that the compiler 404
20 includes optional logic 416 that is configured to reference non-public resources in another
21 assembly providing that these non-public resources have been rendered visible by virtue
22 of an IVT attribute in that other assembly. The IVT attribute can include the optional
23 version argument. Accordingly, for a particular IVT attribute, the logic 416 determines
24 whether the version of the assembly that it is being processed matches the version of the
25 assembly that is identified by the IVT attribute. The logic 416 can perform this function

1 using different matching paradigms. In one case, the logic 416 demands an exact match
2 before an assembly that is being compiled can gain access to the IVT-enabled non-public
3 resources in another assembly. In another case, the logic 416 can be configured to relax
4 this exact matching requirement. For instance, the logic 416 can ignore one or more
5 fields of the version number, or ignore the version number completely. Alternatively, the
6 IVT attribute itself (specified in the referenced assembly) can express the version number
7 by omitting certain fields or providing wildcards for certain fields to relax the matching
8 criteria. In another case, the logic 416 can use a redirection procedure to relax the exact
9 matching requirement (e.g., where a code unit A grants access to some version D, but at
10 compile time, logic 416 redirects from that version of D to a different one, thus also
11 granting access to that new version).

12 For example, assume that assembly A 302 includes an IVT attribute that names a
13 specific version of assembly D 308 as its visibility-privileged assembly. For example,
14 assembly A 302 includes the following IVT attribute naming assembly D 308:

```
15 [assembly:InternalsVisibleToAttribute("D", Version= "v1.0.0.0")]
```

16 When building assembly D 308, the compiler 404 will encounter a reference to types
17 defined as non-public in assembly A 302. As mentioned above, the compiler's logic 416
18 is configured to apply one or more predefined rules that govern the level of exactness that
19 is used in determining whether the actual assembly D 308 satisfies the version argument
20 in the IVT attribute. In one case, the logic 416 requires an exact match, which requires
21 that the assembly D being compiled have precisely version v1.0.0.0. In another case, the
22 logic 416 can apply "fuzzy" matching, such that it will accept any version of assembly D
23 308 that matches the version criteria v1.0.0.*, or version criteria v1.0.*.*, or version
24 criteria v1.*.*.*, and so on (where the symbol "*" represents a wildcard, e.g., a "don't
25 care" number). Alternatively, the logic 416 can be configured to perform a case-by-case

1 investigation to determine the versions of assembly D 308 that the compiler 404 will
2 accept (e.g., by discovering what BindingRedirects apply at compile time). In still
3 another case, the compiler 404 can be configured to entirely ignore the version number
4 specified in the IVT attribute. In this case, the compiler can be configured to grant
5 visibility if other items in the IVT match, such as assembly name, security code, culture,
6 and so on. In this case, version checking can be deferred until runtime.

7 As to runtime, the logic 506 used in the loader/JIT component 506 (of Fig. 5) can,
8 like the compiler 404 (of Fig. 4), apply different version matching rules. In one case, the
9 logic 506 can insist on an exact match between the version being processed by the
10 loader/JIT component 504 and the version identified in the IVT attribute. In other cases,
11 the logic 506 permits various types of “fuzzy” matching rules, where the requirement for
12 exact matching is relaxed. In another case, the logic 506 can use a redirection procedure
13 to relax the exact matching requirement based on policy files (e.g., where a code unit A
14 grants access to some version D, but at runtime, logic 506 redirects from that version of
15 D to a different one, thus also granting access to that new version).

16 For example, assume that assembly D 308 is being compiled which is granted
17 access to non-public resources in assembly A 302 by virtue of the IVT attribute added to
18 assembly A 302 identified above (which specifies version v1.0.0.0). A publisher policy
19 associated with assembly D 308 may state that any reference to versions v1.0.0.0 through
20 v1.1.99.99 should be redirected to version 1.2.0.0. This can be implemented by the
21 following Extensible Markup Language (XML) statements:

```
22  
23 <assemblyIdentity name="D" publicKeyToken="32ab4ba45e0a69a1" />  
24 <bindingRedirect oldVersion="1.0.0.0-1.1.99.99" newVersion="1.2.0.0" />  
25
```

1 In this case, at runtime, the loader/JIT component 504 loads and makes a reference to a
2 non-public type in assembly A 302. Then, the loader/JIT component 504 notes the
3 version of assembly D 308 from which this particular reference is being made and
4 determines whether assembly A 302's IVT attribute-specified version of assembly D 308
5 (i.e., version "v1.0.0.0") is being redirected to the requesting version of assembly D. If
6 this is the case, then the visibility check succeeds. For example, in the BindingRedirect
7 example above, a reference from assembly D version "1.2.0.0" would succeed, because it
8 is defined as a substitutable visibility-privileged assembly for assembly A 302's original
9 declared visibility-privileged assembly version "v1.0.0.0."

10 *A.2. The Security-Related Argument*

11 As noted above, an assembly can be identified in an IVT attribute by specifying
12 its name, as well as any combination of the optional parameters of version, culture, and
13 security code (or perhaps none of these optional parameters). Various security codes and
14 corresponding mechanisms can be employed depending on the requirements of a
15 particular application; some security codes may be more secure (e.g., more robust or
16 stronger) than others. For example, the security code can be implemented as a public key
17 token of any size, a full public key, or some other way of ensuring identity. A named
18 assembly that is accompanied by a security code that ensures identity is referred to as a
19 "strong-named" assembly. Alternatively, the security code can be implemented as any
20 kind of "loose" matching mechanism that provides some security assurances but cannot
21 guarantee identity.

22 More specifically, in one exemplary implementation, the security code can be
23 implemented as a public key token (e.g., PublicKeyToken in the .NET programming
24 environment) that is formed as an 8-byte or 16-byte cryptographically strong hash of a
25 full 160-byte public key (or any other byte-size hash).

1 In one implementation, the developer can specify the public key token by
2 discovering the public key token of the visibility-privileged assembly, and then typing
3 this public key token as an argument of an IVT attribute within the source code that will
4 produce the assembly. For example, assuming that the public key token is
5 "32ab4ba45e0a69a1," then the developer can form the following IVT attribute at an
6 appropriate location within the assembly:

```
7  
8 [assembly:VisibleToAttribute("D", PublicKeyToken= "32ab4ba45e0a69a1")]  
9
```

10 Alternatively, the developer can specify IVT attributes within the source code which lack
11 one or more arguments, such as the public key tokens. In this case, the IVT attribute
12 would appear as follows in the source code:

```
13  
14 [assembly:VisibleToAttribute("D")]  
15
```

16 The missing public key token can be supplied via the command line as follows:

```
17  
18 csc MyApp.cs /keyfile:D:bpub  
19
```

20 This command line information directs the compiler 404 to look in keyfile "bpub" to find
21 the public key token that assembly D contains (or, at least, will contain, after it is
22 eventually built).

23 *A.3. Effect of the IVT Attribute on Type Members*

24 Again assume that an IVT attribute is provided in a compiland that becomes part
25 of an assembly called A 302, and that this IVT attribute designates assembly D 308 as a

1 visibility-accessible assembly (that is, $A \leq D$). This declaration means that assembly A
2 302's types are visible to assembly D 308. Thus, whilst compiling any file that goes to
3 make up assembly D, the compiler will regard all types defined within assembly A as
4 visible. In particular, all "top level" types declared in assembly A 302 become visible
5 (where the public types were already visible to all assemblies by default; hence, the IVT
6 attribute effectively augments the public types by adding the non-public types for
7 visibility-privileged assembly D 308). However, even though a type is visible to a piece
8 of code, its members are not necessarily accessible (in one exemplary implementation).
9 (Here, "members" refers to any field, property, event, or method of the type.) More
10 specifically, in one implementation, various accessibility rules may restrict the
11 accessibility of members that are nonetheless rendered visible by an IVT attribute.

12 For example, consider the following class "Foo" identified in Example 1,
13 expressed in both pseudo CLR code and corresponding exemplary C# code:
14
15
16
17
18
19
20
21
22
23
24
25

Example 1: Type Member Accessibility

| Pseudo-code, using CLR names | C# code, where applicable |
|------------------------------|-----------------------------------|
| non-public class Foo | class Foo { // default = internal |
| PrivateScope a | // not expressible in C# |
| Private b | private int b; |
| FamANDAssem c | // not expressible in C# |
| Assem d | internal int d; |
| Family e | protected int e; |
| FamORAssem f | protected internal int f; |
| Public g | public int g; |
| Endclass | } |

The terms “PrivateScope,” “Private,” “FamANDAssem,” “Assem,” “Family,” and “FamORAssem” refer to different accessibility rules applicable to the exemplary environment of .NET Framework members or other code environment members. Namely, the term “PrivateScope” effectively means that the member is visible within a module (subject to various technology-dependent qualifications). The term “Private” specifies that the member is accessible only from within the same type as the member or within a nested type. The term “Family” specifies that the member is accessible from within the same type as the member and from derived types that inherit from it. The term “Assem” specifies that the member is accessible only in the assembly in which the type is defined. The term “FamANDAssem” means that the member is accessible only from types that qualify for both family and assembly access. The term “FamORAssem” means

1 that the assembly is accessible only from types that qualify for either family or assembly
2 access. And the term “Public” means that the type is accessible from any type.

3 For the example specified above, even though an IVT attribute may have
4 specified that the class Foo is visible to the assembly D 308, assembly D 308 cannot
5 access members “a” or “b” (because these members are designated as private). In other
6 words, the visibility-privilege relationship does not override the private designation on
7 members.

8 As to member c, code within sub-classes of Foo, in assembly D 308, can access
9 this member. In effect, the visibility-privileged relationship has widened “Assem” to
10 span assembly D 308. Normally, it is not possible to define sub-classes of Foo outside its
11 defining assembly; but the visibility-privileged relationship allows a developer to do so
12 within assembly D 308.

13 As to member d, all code in assembly D 308 can access this member. In effect,
14 the visibility-privileged relationship has widened “Assem” to span assembly D 308.

15 As to member e, code within sub-classes of Foo, in assembly D 308, can access
16 this member. That is, there is no change due to the visibility-privileged relationship.

17 As to member f, all code in assembly D 308 can access this member. In effect,
18 the visibility-privileged relationship has widened “Assem” to span assembly D 308.

19 As to member g, all code within assembly D 308 can access this member. That is,
20 there is no change due to visibility-privileged relationship.

21 Generally, note that any CLR accessibility that includes “Assem” is affected by
22 the visibility-privileged relationship. In effect, the definition of “assembly” accessibility
23 is widened to include the visibility-privileged assembly (but strictly in the one-way sense
24 of applying to type members defined in assembly A 302). Similarly, the “family”
25 accessibility becomes widened to visibility-privileged assemblies too. More specifically,

for static members, the general rule is that a family member can be accessed by code within the type that defines that member, or in any sub-type; for the case of instance members, the rule is that family members can be accessed by code in a sub-type so long as this access is performed via an object reference that is of that caller's sub-type (or one of its sub-types).

A.4. The Case of Nested Types

The above analysis applies to so-called "top-level" member types (that is, non-nested member types). In the case of non-top-level (i.e., nested) member types, their accessibility can be narrowed by the accessibility of their enclosing type. For example, consider the following example in which a class X encloses a subclass XX, which, in turn, encloses another subclass XXX. Further assume that the outer enclosing class X is designed as non-public, but has otherwise been rendered visible to assembly D 308.

Example 2: Nested Member Type Accessibility

| Pseudo-code, using CLR names | C# code, where applicable |
|------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <pre> non-public class X public class XX Family class XXX Public e endclass endclass endclass </pre> | <pre> class X { // public class XX { protected class XXX { public int e; } } } </pre> |

1 In this case, the field “e” within subclass XXX is public. But this field’s
2 containing class, XXX, has “Family” accessibility within its enclosing class, XX. So, the
3 only code that can access field “e” is code within class XXX itself, or in a class derived
4 from XX.

5 However, there is an additional complication that confronts compilers in the
6 above example. If there is only one compiland that goes to make up assembly D 308,
7 then the compiler 404 has no difficulty performing the tasks described above, that is, in
8 checking whether assembly D 408 can access types in assembly A 302. However, if the
9 compiler 404 is compiling separate files that eventually merge into the same assembly, it
10 may lack sufficient information for performing this merging task. Assume, for example,
11 that the compiler 404 is separately compiling three files labeled D1, D2 and D3 into
12 separate modules. At compilation, the compiler 404 may not know what assembly they
13 will become bound to. In one implementation, this problem can be addressed by defining
14 a new custom attribute, the “PartOfAssembly” attribute (again, this attribute name is
15 descriptive of its function, but this term is otherwise exemplary and non-limiting). A
16 developer can add this attribute into each compiland to provide the required reference
17 information. For example, this attribute can be expressed using the exemplary syntax of:

18
19 `[module:PartOfAssemblyAttribute("D")]`
20

21 This attribute specifies that the module that receives this attribute is a part of assembly D
22 308.

23 Alternatively, the “part of” reference information can be supplied via the
24 command line in the manner discussed above. For example, the following command line
25

1 information can be used to instruct the compiler 404 that modules D1, D2, and D3 should
2 be compiled into assembly D 308.

3
4 csc /target:module /assembly:D D1.cs

5 csc /target:module /assembly:D D2.cs

6 csc /target:module /assembly:D D3.cs
7

8 (In the specific context of a .NET programming environment, reference to the eventual
9 “home” assembly should be preserved into that copy of the .netmodule, so that ALink can
10 check that the developer is not including the module into a different, non-visibility-
11 privileged assembly; otherwise, this might represent a security risk.)

12 *A.5 Impact on Reflection*

13 A reflection operation allows a program to interact with metadata associated with
14 an assembly, such as during runtime. For instance, the reflection operation can be used to
15 view metadata, discover types, and so on. A reflection emit procedure enables the
16 creation of new types during runtime in a dynamic fashion.

17 In one implementation, the reflection functionality is not modified to
18 accommodate the inclusion of IVT functionality. That is, a developer can continue to
19 successfully query whether a type is public or non-public, but, in this exemplary
20 implementation, the developer cannot query whether a type is visible from another type
21 by virtue of an IVT attribute; however, in another implementation, logic can be provided
22 to accommodate this supplemental querying function. In any case, reflection can be used
23 to enumerate any IVT attributes within any assembly.
24
25

1 In one implementation, no changes are made to reflection emit to support
2 visibility-privileged assemblies. A user can use current reflection emit methods to build a
3 custom attribute to denote any required visibility-privileged relationship.

4 5 **E. Exemplary Computer Environment**

6 Fig. 7 provides information regarding a computer environment 700 that can be
7 used to implement any of the processing functions described in the proceeding sections,
8 such as various compilation operations provided by the source code compiler 404 (of Fig.
9 4) and/or various execution operations provided by the loader/JIT component 504 (of Fig.
10 5).

11 The computing environment 700 includes the general purpose computer 702 and
12 the display device 704 discussed in the context of Fig. 1. However, the computing
13 environment 700 can include other kinds of computer and network architectures. For
14 example, although not shown, the computer environment 700 can include hand-held or
15 laptop devices, set top boxes, programmable consumer electronics, mainframe
16 computers, gaming consoles, etc. Further, Fig. 7 shows elements of the computer
17 environment 700 grouped together to facilitate discussion. However, the computing
18 environment 700 can employ a distributed processing configuration. In a distributed
19 computing environment, computing resources can be physically dispersed throughout the
20 environment.

21 Exemplary computer 702 includes one or more processors or processing units
22 706, a system memory 708, and a bus 710. The bus 710 connects various system
23 components together. For instance, the bus 710 connects the processor 706 to the system
24 memory 708. The bus 710 can be implemented using any kind of bus structure or
25 combination of bus structures, including a memory bus or memory controller, a

1 peripheral bus, an accelerated graphics port, and a processor or local bus using any of a
2 variety of bus architectures. For example, such architectures can include an Industry
3 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced
4 ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a
5 Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

6 Computer 702 can also include a variety of computer readable media, including a
7 variety of types of volatile and non-volatile media, each of which can be removable or
8 non-removable. For example, system memory 708 includes computer readable media in
9 the form of volatile memory, such as random access memory (RAM) 712, and non-
10 volatile memory, such as read only memory (ROM) 714. ROM 714 includes an
11 input/output system (BIOS) 716 that contains the basic routines that help to transfer
12 information between elements within computer 702, such as during start-up. RAM 712
13 typically contains data and/or program modules in a form that can be quickly accessed by
14 processing unit 706.

15 Other kinds of computer storage media include a hard disk drive 718 for reading
16 from and writing to a non-removable, non-volatile magnetic media, a magnetic disk drive
17 720 for reading from and writing to a removable, non-volatile magnetic disk 722 (e.g., a
18 "floppy disk"), and an optical disk drive 724 for reading from and/or writing to a
19 removable, non-volatile optical disk 726 such as a CD-ROM, DVD-ROM, or other
20 optical media. The hard disk drive 718, magnetic disk drive 720, and optical disk drive
21 724 are each connected to the system bus 710 by one or more data media interfaces 728.
22 Alternatively, the hard disk drive 718, magnetic disk drive 720, and optical disk drive 724
23 can be connected to the system bus 710 by a SCSI interface (not shown), or other
24 coupling mechanism. Although not shown, the computer 702 can include other types of
25 computer readable media, such as magnetic cassettes or other magnetic storage devices,

1 flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical storage,
2 electrically erasable programmable read-only memory (EEPROM), etc.

3 Generally, the above-identified computer readable media provide non-volatile
4 storage of computer readable instructions, data structures, program modules, and other
5 data for use by computer 702. For instance, the readable media can store the operating
6 system 730, one or more application programs 732, other program modules 734, and
7 program data 736.

8 The computer environment 700 can include a variety of input devices. For
9 instance, the computer environment 700 includes the keyboard 738 and a pointing device
10 740 (e.g., a "mouse") for entering commands and information into computer 702. The
11 computer environment 700 can include other input devices (not illustrated), such as a
12 microphone, joystick, game pad, satellite dish, serial port, scanner, card reading devices,
13 digital or video camera, etc. Input/output interfaces 742 couple the input devices to the
14 processing unit 706. More generally, input devices can be coupled to the computer 702
15 through any kind of interface and bus structures, such as a parallel port, serial port, game
16 port, universal serial bus (USB) port, etc.

17 The computer environment 700 also includes the display device 704. A video
18 adapter 744 couples the display device 704 to the bus 710. In addition to the display
19 device 704, the computer environment 700 can include other output peripheral devices,
20 such as speakers (not shown), a printer (not shown), etc.

21 Computer 702 can operate in a networked environment using logical connections
22 to one or more remote computers, such as a remote computing device 746. The remote
23 computing device 746 can comprise any kind of computer equipment, including a general
24 purpose personal computer, portable computer, a server, a router, a network computer, a
25 peer device or other common network node, etc. Remote computing device 746 can

1 include all of the features discussed above with respect to computer 702, or some subset
2 thereof.

3 Any type of network can be used to couple the computer 702 with remote
4 computing device 746, such as a local area network (LAN) 748, or a wide area network
5 (WAN) 750 (such as the Internet). When implemented in a LAN networking
6 environment, the computer 702 connects to local network 748 via a network interface or
7 adapter 752. When implemented in a WAN networking environment, the computer 702
8 can connect to the WAN 750 via a modem 754 or other connection strategy. The modem
9 754 can be located internal or external to computer 702, and can be connected to the bus
10 710 via serial I/O interfaces 756 or other appropriate coupling mechanism. Although not
11 illustrated, the computing environment 700 can provide wireless communication
12 functionality for connecting computer 702 with remote computing device 746 (e.g., via
13 modulated radio signals, modulated infrared signals, etc.).

14 In a networked environment, the computer 702 can draw from program modules
15 stored in a remote memory storage device 758. Generally, the depiction of program
16 modules as discrete blocks in Fig. 7 serves only to facilitate discussion; in actuality, the
17 programs modules can be distributed over the computing environment 700, and this
18 distribution can change in a dynamic fashion as the modules are executed by the
19 processing unit 706.

20 Wherever physically stored, one or more memory modules 708, 722, 726, 758,
21 etc. can be provided to store the compilation operations described in Fig. 4 and/or the
22 loader/JIT operations described in Fig. 5.

23
24 Although the invention has been described in language specific to structural
25 features and/or methodological acts, it is to be understood that the invention defined in

1 the appended claims is not necessarily limited to the specific features or acts described.
2 Rather, the specific features and acts are disclosed as exemplary forms of implementing
3 the claimed invention.
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25